

CLAIMS

What is claimed is:

1. A system comprising:
 - a first node including data having an associated state, the associated state of the data at the first node being a modified state; and
 - a second node operative to provide a non-migratory source broadcast request for the data, the first node being operative in response to the non-migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an owner state without updating memory, the second node being operative to receive the data from the first node and assign a shared state to an associated state of the data at the second node.
2. The system of claim 1, wherein the second node is further operative to provide a migratory source broadcast request for the data, the first node being operative in response to the migratory source broadcast request to provide the data to the second node and to transition the associated state of the data at the first node from the modified state to an invalid state without updating memory, the second node being operative to receive the data from the first node and assign the associated state of the data at the second node to a dirty state.
3. The system of claim 2, wherein the first node is operative in response to the migratory source broadcast request to provide an ownership data response to the second node.
4. The system of claim 2, wherein the associated state of the data at the second node being the dirty state makes the data at the second node available for migration to other nodes.
5. The system of claim 2, wherein the second node is further operative to write the data to the second node and transition the associated state of the data at the second node from the dirty state to the modified state, the associated state of the data at the second node being the modified state making the data at the second node available for migration to other nodes providing a migratory source broadcast request for the data.

6. The system of claim 2, wherein the second node is programmed with instructions to selectively invoke one of the non-migratory source broadcast request and the migratory source broadcast request.

7. The system of claim 2, wherein the second node is programmed with instructions that provide a predictive selection to invoke one of the non-migratory source broadcast request and the migratory source broadcast request.

8. The system of claim 1, wherein the first node is operative in response to the non-migratory source broadcast request to provide a shared data response to the second node.

9. The system of claim 1, wherein further migration of the data from the second node is precluded when the associated state of the data at the second node is the shared state.

10. The system of claim 1, further comprising at least one other node that provides a non-data response to the second node in response to the non-migratory source broadcast request from the second node, the non-data response indicating that the at least one other node does not have a valid copy of the data requested by the second node.

11. The system of claim 1, wherein the first node defines a first processor and the second node defines a second processor, the first and second processors each having an associated cache that comprises a plurality of cache lines, each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line, the first and second processors being capable of communicating with each other and with a system memory via an interconnect, the system further comprising a first cache controller associated with the first processor and a second cache controller associated with the second processor, the first cache controller being operative to manage data requests and responses for the associated cache of the first processor, the first cache controller effecting state transitions associated with the data in the associated cache of the first processor based on the data requests and responses for the associated cache of the first processor, the second cache controller being operative to manage data requests and responses for the associated cache of the second processor, the second cache controller effecting state

transitions associated with the data in the associated cache of the second processor based on the data requests and responses for the associated cache of the second processor.

12. A multi-processor network comprising:

memory for storing data;

a first processor node having a first processor node cache line including the data, the first processor node cache line having an associated state, the associated state of the first processor node cache line being a modified state; and

a second processor node operative to provide a non-migratory source broadcast read request for the data, the second processor node having a second processor node cache line with an associated state;

the first processor node being programmed to respond to the non-migratory source broadcast read request of the second processor node by providing a shared data response to the second processor node and transitioning the associated state of the first processor node cache line from the modified state to an owner state without updating the memory with the data, the data being stored in the second processor node cache line, the associated state of the second processor node cache line being assigned a shared state.

13. The network of claim 12, wherein the second processor node is further operative to provide a migratory source broadcast read request for the data, the first processor node being programmed to respond to the migratory source broadcast read request of the second processor node by providing an ownership data response to the second processor node and transitioning the associated state of the first processor node cache line from the modified state to an invalid state without updating the memory with the data, the data from the ownership data response being stored in the second processor node cache line and the state associated with the second processor node cache line being assigned a dirty state.

14. The network of claim 13, wherein the data stored in the second processor node cache line assigned the dirty state is available for migration to other nodes.

15. The network of claim 13, wherein the second processor node is further operative to write the data stored in the second processor node cache line assigned the dirty state and transition the state associated with the second processor node cache line from the dirty state to the modified state, the data stored in the second processor node cache line assigned

the modified state being available for migration to other nodes providing a migratory source broadcast read request for the data.

16. The network of claim 13, wherein the second processor node is programmed with instructions to selectively invoke one of the non-migratory source broadcast read request and the migratory source broadcast read request to obtain the data.

17. The network of claim 13, wherein the second processor node is programmed with instructions to predictively select one of the non-migratory source broadcast read request and the migratory source broadcast read request to obtain the data.

18. The network of claim 12, wherein further migration of the data from the second processor node is precluded when in the shared state.

19. The network of claim 12, further comprising at least one other processor node that provides a non-data response to the second processor node in response to the non-migratory source broadcast read request from the second processor node, the non-data response indicating that the at least one other processor node does not have a valid copy of the data requested by the second processor node.

20. The network of claim 12, wherein the first and second processor nodes each have an associated cache that comprises a plurality of cache lines, each cache line having a respective tag address that identifies associated data and each cache line having state information that indicates a state of the associated data for the respective cache line, the first and second processor nodes being capable of communicating with each other and with the memory via an interconnect, the system further comprising a first cache controller associated with the first processor node and a second cache controller associated with the second processor node, the first cache controller being operative to manage data requests and responses for the associated cache of the first processor node, the first cache controller effecting state transitions associated with the data in the cache of the first processor node based on the data requests and responses for the associated cache of the first processor node, the second cache controller being operative to manage data requests and responses for the associated cache of the second processor node, the second cache controller effecting state transitions associated with the data in the associated cache of the second

processor node based on the data requests and responses for the associated cache of the second processor node.

21. A computer system comprising:

a source processor having an associated source processor cache, the source processor being operative to issue a selected one of a non-migratory source broadcast (XREADN) request for data and a migratory source broadcast (XREADM) request for data;

memory storing the data; and

a target processor having an associated target processor cache with a target processor cache line that stores the data, the target processor cache line having an associated state, the associated state of the target processor cache line being a modified state, the target processor being programmed to respond to the XREADN request by providing a shared data (S-DATA) response to the source processor and by transitioning the associated state of the target processor cache line from the modified state to an owner state without updating the memory, the target processor being programmed to respond to the XREADM request by providing an ownership data (D-DATA) response to the source processor and by transitioning the associated state of the target processor cache line from the modified state to an invalid state without updating the memory.

22. The computer system of claim 21, wherein the source processor further comprises an associated source processor cache having a source processor cache line for storing the data, the source processor cache line having an associated state, the source processor storing the data in the source processor cache line and assigning a shared state to the associated state of the source processor cache line in response to receiving the S-DATA response from the target processor.

23. The computer system of claim 21, wherein the source processor further comprises an associated source processor cache having a source processor cache line for storing the data, the source processor cache line having an associated state, the source processor storing the data in the source processor cache line and assigning a dirty state to the associated state of the source processor cache line in response to receiving the D-DATA response from the target processor.

24. A system comprising:

means for broadcasting from a first node a non-migratory read (XREADN) request for data;

means for providing the data from a second node to the first node in response to the XREADN request, a modified state being associated with the data at the second node, a shared state being associated with the data at the first node in response to the first node receiving the data from the second node; and

means for transitioning the modified state associated with the data at the second node to an owner state without updating memory of the system.

25. The system of claim 24, further comprising:

means for broadcasting from the first node a migratory read (XREADM) request for data;

means for providing the data from the second node to the first node in response to the XREADM request, the modified state being associated with the data at the second node, the shared state being associated with the data at the first node in response to the first node receiving the data from the second node; and

means for transitioning the modified state associated with the data at the second node to an invalid state without updating memory of the system.

26. The system of claim 24, further comprising means for selecting one of the XREADM request and XREADN request to broadcast from the first node.

27. The system of claim 24, further comprising means for predictively selecting one of the XREADM request and XREADN request to broadcast from the first node.

28. A method comprising:

broadcasting a non-migratory request for data from a first node to other nodes of an associated system;

providing a shared copy of the data from a second node to the first node in response to the non-migratory request;

transitioning a state associated with the data at the second node from a modified state to an owner state in response to the non-migratory request; and

transitioning a state associated with the data at the first node to a shared state in response to receiving the shared copy of the data from the second node.

29. The method of claim 28, further comprising:

broadcasting a migratory request for the data from the first node to other nodes of the associated system;

providing an ownership data response from the second node to the first node in response to the migratory request;

transitioning the state associated with the data at the second node from a modified state to an invalid state data in response to the migratory request; and

transitioning the state associated with the data at the first node to a dirty state in response to receiving the ownership data response from the second node.

30. The method of claim 29, further comprising selecting one of the migratory request and the non-migratory request to broadcast from the first node.

31. The method of claim 29, further comprising predictively selecting one of the migratory request and the non-migratory request to broadcast from the first node.

32. A computer system comprising a cache coherency protocol that is operative to permit migration of data to a cache associated with a source processor from a cache associated with a target processor when a migratory request is issued from the source processor, the protocol being further operative to prevent migration of the data to the cache associated with the source processor from the cache associated with the target processor when a non-migratory request is issued from the source processor.